

EXTERNAL STORAGE SUBSYSTEM

BACKGROUND OF THE INVENTION

1. The present invention relates to an external storage subsystem, and more particularly to a technology effective for an improvement of reliability of an external storage subsystem having a cache function.
2. For example, in a magnetic disk subsystem used as an external storage in a general purpose computer system, a well-known cache memory comprising a semiconductor memory is interleaved at a portion of a disk controller to avoid as much as possible the reduction of a data transfer rate due to a mechanical factor such as a rotational delay time or a latency time in a magnetic disk drive.
3. A cache structure in such a disk controller is discussed in "A Multiport Page-Memory Architecture and A Multiport Disk-Cache System" New Generation Computing 2 (1984) 241-260 OHMSHA, LTD. and Springer-Verlag, in which it is proposed to improve an access performance to the cache by dividing into a plurality of memory banks. Further, a switching network called an interconnection network is proposed as a system for coupling the memory banks and a channel or a disk controller.
4. The conventional technology above intends to improve the cache performance by providing a plurality of memory banks and the switching network. As to a data

bus structure in the disk controller, the switching network system called the interconnection network is proposed. However, the switching network system is imparted with a hardware restriction when a data bus configuration for exchanging data is to be constructed by connecting a plurality of memory banks and a plurality of channel units or a plurality of control units.

1. It does not refer to the multiplexity of the cache unit comprising the memory banks.

SUMMARY OF THE INVENTION

2. It is an object of the present invention to provide a data bus structure for connecting a plurality of cache units of a host and a plurality of channel units or a plurality of control units of a rotating storage, by taking restrictive conditions of a data transfer rate and a data bus width on hardware into consideration.

3. It is another object of the present invention to provide an external memory subsystem which has a high tolerance to failures and has a highly reliable cache function.

4. The above and other objects and features of the present invention will be apparent from the following description of the present invention when taken in conjunction with the attached drawings.

1. Representative features of the present invention are briefly explained below.
2. The external storage subsystem of the present invention comprises a rotating storage for storing data to be accessed from a host and an external memory control unit having a cache mechanism for responding to an access request from the host to the rotating storage by temporarily holding data exchanged between the rotating storage and the host and having at least one of a non-volatile (persistent) semiconductor memory and a volatile (non-persistent) semiconductor memory as a storage medium. It further comprises a plurality of independent cache mechanisms and a plurality of independent access paths for permitting independent accesses from the host and the rotating storage to the respective cache mechanisms.
3. In the external storage subsystem of the present invention, the external memory control unit comprises a plurality of channel units for controlling the transfer of data to and from the host and a plurality of control units for controlling the transfer of data to and from the rotating storage, and each of the channel units and the control units has a plurality of the first access paths to which a plurality of cache mechanisms are to be independently connected.
4. In the external storage subsystem of the present invention, the external memory control unit comprises a plurality of channel units for controlling the transfer of data to and from the host and a

plurality of control units for controlling the transfer of data to and from the rotating storage, and each of the cache mechanisms has a plurality of the second access paths to which the channel units and the control units are to be connected.

1. In the external storage subsystem of the present invention, the external memory control unit comprises a plurality of channel units for controlling the transfer of data to and from the host, a plurality of control units for controlling the transfer of data to and from the rotating storage, and a plurality of independent access paths to the channel units, the control units and the cache mechanisms. The respective channel units, control units and cache mechanisms are connected to the third access paths.

2. In the external storage subsystem of the present invention, the external memory control unit comprises a plurality of channel units for controlling the transfer of data to and from the host, a plurality of control units for controlling the transfer of data to and from the rotating storage, and the fourth access paths for directly and independently connecting the respective channel units and control units with the respective cache mechanisms.

3. In the external storage subsystem of the present invention, since the cache units are multiplexed and the access paths to the respective cache units by the host and the rotating storage are of independent

configuration, the data transfer rate or the data bus width can be optimized by combining a plurality of cache units and a plurality of channel units or a plurality of control units.

1. Further, since the cache units and the access paths to the cache units are multiplexed, a probability of maintaining the cache function in case a trouble occurs is enhanced and the reliability of the external storage subsystem and the tolerance to the failures are certainly improved.

2. The effects of the representative features of the present invention are as follows.

3. In the external storage subsystem of the present invention, the cache units in the external storage subsystem including the rotating storage can be coupled, in a simple construction, to the channel units of the host and the control units of the rotating storage. Accordingly, the cache function and performance in the disk control unit are improved.

4. Further, in the external storage subsystem of the present invention, since both the cache units and the access paths to the cache units are multiplexed, the tolerance to the failures is high and the highly reliable cache function is attained.

BRIEF DESCRIPTION OF THE DRAWINGS

5. Fig. 1 shows a block diagram of one embodiment of an external storage subsystem of the present invention,

1. Fig. 2 shows a signal configuration of a data bus connecting a channel unit or control unit in a disk controller and a plurality of cache units or nonvolatile memory units,
2. Fig. 3 shows a conceptual view of data bus protocol when read data, write data and command status are exchanged between the channel unit or the control unit and the cache units or the nonvolatile memory units,
3. Fig. 4 illustrates a data bus mode for specifying a status of a data bus,
4. Fig. 5 shows a block diagram of a configuration of another embodiment of the external storage subsystem of the present invention,
5. Fig. 6 shows a block diagram of other embodiment of the external storage subsystem of the present invention, and
6. Fig. 7 shows a block diagram of a further embodiment of the external storage subsystem of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[Embodiment 1]

7. One embodiment of the external memory subsystem of the present invention is explained with reference to the drawings here.

1. As shown in Fig. 1, a computer system of the present embodiment comprises a central processing unit (CPU) 1 and a disk subsystem which includes a disk controller 2 and a magnetic disk drive 3.

2. The CPU 1 and the disk controller 2 are connected through a plurality of channel interfaces 4, and the disk controller 2 and the magnetic disk drive 3 are connected through a plurality of control interfaces 5.

3. The CPU 1 issues an access command to the disk controller 2 through the channel interface 4 to control the read and write of data by the magnetic disk drive 3 by the disk controller 2 through the control interface 5 in accordance with the command from the CPU 1. In other expression, the CPU 1 controls the read and write of data by controlling the magnetic disk drive 3 by the disk controller 2 through the control interface 5 in accordance with the access command issued to the disk controller 2 through the channel interface 4 from the CPU 1.

4. In the disk controller 2, a plurality of channel units 60 and channel units 61 which operate under a channel control processor 110 and a channel control processor 111 are provided on the side of the channel interface 4, and control units 70 and control units 71 which operate under a control unit control processor 120 and a control unit control processor 121 are provided on the side of the control interface 5.

5. A command issued from the CPU 1 to the disk

controller 2 is accepted by the channel units 60 and 61, decoded by the channel control processors 110 and 111 and delivered to the control unit processors 120 and 121 necessary for controlling the magnetic disk drive 3. The control unit processors 120 and 121 control the magnetic disk drive 3 through the control units 70 and 71.

1. The disk controller 2 includes two independent cache units 80 and 81 for temporarily storing data in a semiconductor memory as a storage medium, which is not described in figure however, and two independent nonvolatile memory units 90 and 91. The nonvolatile memory units 90 and 91 are rewritable memories and have a capability of holding data for a certain time period without regard to the presence or absence of external power supply.
2. Each capacity of the cache units 80 and 81, for example, is determined to permit a sufficient individual cache operation compatible to a memory capacity of the magnetic disk drive 3. Similarly, a capacity of each of the nonvolatile memory units 90 and 91, for example, is determined to permit a sufficient individual cache operation compatible to the memory capacity of the magnetic disk drive 3.
3. The channel unit 60 is provided with a plurality of independent data buses 60A and 60B. The cache unit 80 and the nonvolatile memory unit 91 are independently connected to the data bus 60A through the

respective access lines 80a and 91a. The cache unit 81 and the nonvolatile memory unit 90 are connected to the data bus 60B through the respective access lines 81a and 90a.

1. Similarly, the channel unit 61 is provided with a plurality of independent data buses 61A and 61B. The cache unit 80 is connected to the data bus 61A through an access line 80c, and the nonvolatile memory unit 91 is connected through an access line 91c. The cache unit 81 is connected to the data bus 61B through an access line 81c and the nonvolatile memory unit 90 is connected through an access line 90c.

2. The control unit 70 is provided with a plurality of independent data buses 70A and 70B. The cache unit 80 is connected to the data bus 70A through an access line 80b, and the nonvolatile memory unit 91 is connected through an access line 91b. The cache unit 81 is connected through an access line 81b and the nonvolatile memory unit 90 is connected through an access line 90b.

3. Similarly, the control unit 71 is provided with a plurality of independent data buses 71A and 71B. The cache unit 80 is connected to the data bus 71A through an access line 80d and the nonvolatile memory unit 91 is connected through an access line 91d. The cache unit 81 is connected to the data bus 71B through an access line 81d and the nonvolatile memory unit 90 is connected through an access line 90d.

1. In the configuration of the present embodiment, each of the channel units 60 and 61 and the control units 70 and 71 can access to the cache units 80 and 81 and the nonvolatile memory units 90 and 91 through independent paths.
2. An operation of the external memory subsystem of the present embodiment is explained below.
3. Write data sent from the CPU 1 to the disk controller 2 is temporarily stored in one of the cache unit 80 and 81 and the nonvolatile memory unit 90 and 91 through the channel unit 60 or 61 and the one of the data buses 60A, 60B, 61A and 61B, in accordance with the command from the channel control processor 110 or 111. Then, the data is read from one of the cache units 80 and 81 or one of the nonvolatile memory units 90 and 91 by the command from the control unit control processor 120 or 121 and the write data is stored in the magnetic disk drive 3 through the data bus between 70A and 71B (70A to 71B) and the control unit 70 or 71.
4. On the other hand, when the channel control processor 110 or 111 receives a data read request from the CPU 1 through the channel unit 60 or 61, it searches the contents of the cache unit 80 or 81 and the nonvolatile memory unit 90 or 91, and if there is data requested by the CPU 1, it sends the read data from the cache unit 80 or 81 or the nonvolatile memory unit 90 or 91 through one of the data bus between 60A and 61B (60A to 61B) and the channel unit 60 or 61.

1. If the data requested by the CPU 1 is not present in any of the cache units 80 and 81 and the nonvolatile memory units 90 and 91, the channel control processor 110 or 111 conveys the data read from the magnetic disk drive 3 to the control unit control processor 120 or 121. When the control unit control processor 120 or 121 receives the data read request from the channel control processor 110 or 111, it stores the requested data read from the magnetic disk drive 3 into one of the cache unit 80 or 81 through the control unit 70 or 71 and the data bus between 70A and 71B (70A to 71B).
2. When the channel control processor 110 or 111 receives a report of completion of read data store into the cache unit 80 or 81 from the control unit control processor 120 or 121, it reports the completion of data to the CPU 1 and sends the read data from the cache unit 80 or 81 to the CPU 1 through the data bus between 60A and 61B (60A to 61B) and the channel unit 60 or 61 in response to the command from the CPU 1.
3. Fig. 2 shows an example of signal configuration of the data buses from 60A through 71B connecting the channel units 60 and 61 or the control units 70 and 71 and the cache units 80 and 81 or the nonvolatile memory units 90 and 91 in the disk controller 2. In the present embodiment, the channel units 60 and 61 or the control units 70 and 71 perform a master operation to the cache units 80 and 81 and the nonvolatile memory

units 90 and 91. The cache units 80 and 81 or the nonvolatile memory units 90 and 91 perform a slave operation to the channel units 60 and 61 or the control units 70 and 71.

1. The channel units 60 and 61 or the control units 70 and 71 select the cache unit 80 or 81 or the nonvolatile memory unit 90 or 91 by driving a SEL (0 - 1) signal line. The channel unit 60 or 61 or the control unit 70 or 71 specifies a status of the data buses from 60A through 71B, by a combination of the signals shown in Fig. 4, of a DTOUT/*DTIN signal line and a CMD/*DTIN signal line in the selected condition of the cache units 80 and 81 or the nonvolatile memory units 90 and 91. The read data, write data, command and status are exchanged between the channel units 60 and 61 or the control units 70 and 71 and the cache units 80 and 81 or the nonvolatile memory units 90 or 91, in accordance with the data bus protocol shown in Fig. 3.

2. In the external storage subsystem of the present embodiment, the cache units 80 and 81 and the nonvolatile memory units 90 and 91 are independently provided, and the accesses to the cache units 80 and 81 and the nonvolatile memory units 90 and 91 from the channel units 60 and 61 on the side of the CPU 1 and the control units 70 and 71 of the magnetic disk device 3 are permitted independently through the data buses 60A - 71B, the access lines 80a to 80d, 81a to 81d, 90a to 90d and 91a - 91d. Accordingly, the data transfer rate and

the data bus width in the data bus or the access line can be optimally set.

1. Further, even if a failure occurs in any of the cache units 80 and 81 and the nonvolatile memory units 90 and 91 or the data buses 60A - 71B or further the access lines 80a - 80d, 81a - 81d, 90a - 90d and 91a - 91d, the cache function can be maintained and the tolerance to the failures and the reliability of the operation are improved.

[Embodiment 2]

2. Fig. 5 shows a block diagram of another embodiment of the external storage subsystem of the present invention.

3. In the embodiment 2, each of the cache units 80 and 81 and the nonvolatile memory units 90 and 91 has a plurality of data buses 80A, 80B, 81A, 81B, 90A, 90B, 91A and 91B to each of which the channel units 60 and 61 and the control units 70 and 71 are independently connected through the access lines 60a to 60d, 61a to 61d, 70a to 70d and 71a to 71d.

4. The similar effects to those of the embodiment 1 are attained in the present embodiment.

[Embodiment 3]

5. Fig. 6 shows a block diagram of a configuration of other embodiment of the external memory subsystem of the present invention.

6. In the embodiment 3, two common data buses 200A and 200B are provided. To each of them, the

channel unit 60 is connected through access lines 60e and 60f, the channel unit 61 is connected through access lines 61e and 61f, the control unit 70 is connected through access lines 70e and 70f, the control unit 71 is connected through access lines 71e and 71f, the cache unit 80 is connected through access lines 80e and 80f, the cache unit 81 is connected through access lines 81e and 81f, the nonvolatile memory unit 90 is connected through access lines 90e and 90f, and the nonvolatile memory unit 91 is connected through access lines 91e and 91f, independently respectively.

1. The similar effects to those of the previous embodiments are attained in the present embodiment.

[Embodiment 4]

2. Fig. 7 shows a block diagram of a configuration of a further embodiment of the external storage subsystem of the present invention.

3. In the embodiment 4, cache unit groups 800 and 801 each comprising a plurality of cache units and nonvolatile memory unit groups 900 and 901 each comprising a plurality of nonvolatile memory units are provided. The cache unit groups 800 and 801 and the nonvolatile memory unit groups 900 and 901 are connected to the channel units 60 and 61 and the control units 70 and 71 through data buses 60g to 60j, 61g to 61j, 70g to 70j and 71g to 71j.

4. The similar effects to those of the previous embodiments are attained in the present embodiment.

1. While the present invention has been described with reference to the embodiments of the present invention, those embodiments are examples of the present invention and they may be modified in various ways without departing from the gist of the invention.